



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/608,037	06/30/2003	Sanjay Ghemawat	0026-0031	8255

44989 7590 07/13/2006

HARRITY SNYDER, LLP
11350 Random Hills Road
SUITE 600
FAIRFAX, VA 22030

EXAMINER

LOVEL, KIMBERLY M

ART UNIT	PAPER NUMBER
----------	--------------

2167

DATE MAILED: 07/13/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

10/608,037

Applicant(s)

GHEMAWAT ET AL.

Examiner

Kimberly Lovel

Art Unit

2167

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 24 April 2006.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-18 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☐ Claim(s) _____ is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 30 June 2003 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

1. This communication is responsive to the Amendment filed 24 April 2006.
2. Claims 1-18 are pending in this application. Claims 1, 13, 14, 15, 16, 17 and 18 are independent. In the Amendment filed 24 April 2006, claim 4 has been amended. This action is made Non-Final.
3. The rejections of claims 1-7 and 12-18 as being anticipated by US Patent No 5,991,804 to Bolosky et al; claims 8 and 9 as being unpatentable over US Patent No. 6,725,392 to Frey et al; and claims 10 and 11 as being unpatentable over US-PGPub 2003/0023898 to Jacobs et al have been withdrawn as necessitated by applicants' arguments.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

4. Claims 1-7 and 12-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent No. 5,991,804 issued to Bolosky et al (hereafter Bolosky et al) in view of US Patent No 6,973,491 to Staveley et al (hereafter Staveley et al).

Referring to claim 1, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality of servers configured to store data (see column 3, lines 47-50 and Fig 1, items 24 and 28 – the data servers are considered to represent the plurality of *servers configured to store data*; item 24 represents the data server which contains item 28 which represents a data disk that stores data) and

a master connected to the servers (see column 3, lines 43-45 and Fig 1, items 22, 24 and 26– the controller is considered to represent the *master*) wherein the master is configured to record location information that identifies ones of the servers that store the data (see column 4, line 64 – column 5, line 3 – the controller can maintain a map that tracks the physical locations of the data blocks).

Bolosky et al fail to explicitly teach the further limitation wherein the master is configured to: communicate with the servers upon startup of the master to authoritatively identify the data stored by the servers and record location information that identifies one of the servers that store the data. Staveley et al disclose a system for monitoring and managing system assets and asset configurations (see abstract). In particular, Staveley et al disclose the further limitation wherein the master is configured to: communicate with the servers (see column 3, lines 50-52 – the master communicates with the devices/machines through the internet; the devices are considered to represent the *servers*) upon startup of the master to authoritatively identify the data stored by the servers (see column 2, lines 59-67 and column 3, line 60 – column 4, line 66 – the program can be executed on a master in a network

Art Unit: 2167

environment; the startup of the program is considered to represent the startup of the master; the program collects data from the devices; the devices are considered to represent the *servers*).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to utilize the data collection programs of Staveley et al as an additional program to the programs of the controller mentioned by Bolosky et al. One would have been motivated to do so since the continuous media file server of Bolosky et al is capable of handling whole system failures and after a whole system failure, the master will need to startup (Bolosky et al: see column 16, lines 7-11).

Referring to claim 2, the combination of Bolosky et al and Staveley et al (hereafter Bolosky/Staveley) discloses the system of claim 1, wherein the data corresponds to files stored as chunks by the servers (Bolosky et al: see column 4, lines 37-44 – the disk stores data in blocks which is considered to represent a chunk since a block is defined as a unit of data or an amount of physical space allocated on a disk to hold that unit of data).

Referring to claim 3, Bolosky/Staveley discloses the system of claim 1, wherein the master is further configured to control placement of new data at the servers (Bolosky et al: see column 8, lines 11-22 – when performing striping, the controller acts as if it were writing the data files for the first time which is considered to represent writing new data).

Referring to claim 4, Bolosky/Staveley discloses the system of claim 3, wherein when controlling the placement of new data, the master is configured to:

identify one or more of the servers to store the new data based on at least one of utilization of the servers (Bolosky et al: column 7, line 64 – column 8, line 2), prior chunk distribution involving the servers (Bolosky et al: see column 8, lines 31-42), network topology, or failure correlation properties associated with the servers, and

place the data at the identified one or more servers (Bolosky et al: see column 8, lines 16-22 – the controller writes the data files to the system).

Referring to claim 5, Bolosky/Staveley discloses the system of claim 1, wherein the master is further configured to control redistribution of the data stored by the servers (Bolosky et al: see column 6, lines 30-38 – the controller is configured to move data blocks around using a process called restripping).

Referring to claim 6, Bolosky/Staveley discloses the system of claim 5, wherein when controlling redistribution of the data, the master is configured to:

select data to redistribute based on a current distribution of the data (Bolosky et al: see column 8, lines 11-22 – the controller selects a location for a first block and then sequentially selects the rest of the blocks),

identify on one or more of the servers to which to move the selected data (Bolosky et al: see column 8, lines 31-53 – the data blocks are moved in parallel), and

move the selected data to the identified one or more servers (Bolosky et al: see column 8, lines 23-30 and column 9, lines 11-43 – the controller begins by placing the blocks on the storage disks; the storage disks are located on the server; the controller moves the data on the bit map).

Referring to claim 7, Bolosky/Staveley discloses the system of claim 1, wherein the master is further configured to monitor the state of the servers (Staveley et al: see column 1, lines 51-58 – the monitoring application monitors the status of the target devices).

Referring to claim 12, Bolosky/Staveley discloses the system of claim 1, wherein the location information is not stored persistently by the master (Bolosky et al: see column 4, line 64 – column 5, line 3 – the system provides two alternatives for the location information; the master does not persistently store the location information in the first alternative mentioned in column 4, line 66 – column 5, line 1).

Referring to claim 13, Bolosky et al disclose in a file system that includes a master connected to a plurality of servers (see abstract and Fig 1), the master comprising: means for storing location information that identifies ones of the servers that stores the data (see column 5, lines 1-3 – the map represents the means).

Bolosky et al fail to explicitly teach the further limitations of means for performing a startup operation (see column 2, lines 55-57 – the controller represents the means; column 7, lines 4-19; and column 16, lines 34-48) and means for communicating with the servers during or after the startup operation to authoritatively identify the data stored by the servers (see column 3, lines 43-45 and Fig 1, item 26 – the low bandwidth control network represents the means). Staveley et al disclose a system for monitoring and managing system assets and asset configurations (see abstract). In particular, Staveley et al disclose means for performing a startup operation (see column 2, lines 59-67 and column 3, line 60 – column 4, line 66 – the configuration of the master provides the

means); and means for communicating with the servers during or after the startup operation to authoritatively identify the data stored by the servers (see column 3, lines 50-52; column 2, lines 59-67; and column 3, line 60 – column 4, line 66 – the program can be executed on a master in a network environment; the startup of the program is considered to represent the startup of the master; the program collects data from the devices; the devices are considered to represent the servers; the master communicates with the devices/machines through the internet).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to utilize the data collection programs of Staveley et al as an additional program to the programs of the controller mentioned by Bolosky et al. One would have been motivated to do so since the continuous media file server of Bolosky et al is capable of handling whole system failures and a after a whole system failure, the master will need to startup (Bolosky et al: see column 16, lines 7-11).

Referring to claim 14, Bolosky et al disclose a method for maintaining data in a file system that includes a master connected to a plurality of servers, the method, performed by the master (see Fig 1), comprising: generating location information based on the data determined to be stored by the servers (see column 4, line 64 – column 5, line 3 – the information is generated on for map, which tracks the location of the data block which identifies which server the data is stored on).

Bolosky et al fail to explicitly teach the further limitation of communicating with the servers upon startup of the master to authoritatively determine data stored by the servers. Staveley et al disclose a method for monitoring and managing system assets

Art Unit: 2167

and asset configurations (see abstract). In particular, Staveley et al disclose the further limitation of communicating with the servers (see column 3, lines 50-52 – the master communicates with the devices/machines through the internet; the devices are considered to represent the *servers*) upon startup of the master to authoritatively determine data stored by the servers (see column 2, lines 59-67 and column 3, line 60 – column 4, line 66 – the program can be executed on a master in a network environment; the startup of the program is considered to represent the startup of the master; the program collects data from the devices; the devices are considered to represent the *servers*).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to utilize the data collection programs of Staveley et al as an additional program to the programs of the controller mentioned by Bolosky et al. One would have been motivated to do so since the continuous media file server of Bolosky et al is capable of handling whole system failures and a after a whole system failure, the master will need to startup (Bolosky et al: see column 16, lines 7-11).

Referring to claim 15, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality of servers configured to store files as chunks (see column 1, line 66 – column 2, line 3 and column 4, lines 37-44 – the disk stores data in blocks which is considered to represent a chunk since a block is defined as a unit of data or an amount of physical space allocated on a disk to hold that unit of data); and

a master connected to the servers (see column 3, lines 43-45 and Fig 1, items 22, 24 and 26 – the controller is considered to represent the *master*).

Bolosky et al fail to explicitly teach the further limitations of authoritatively determining location information by communicating with the servers, the location information being based on which of the servers stores ones of the chunks and updating the location information by periodically communicating with the servers to obtain changes to the location information. Staveley et al disclose a system for monitoring and managing system assets and asset configurations (see abstract). In particular, Staveley et al disclose authoritatively determining location information by communicating with the servers, the location information being based on which of the servers stores ones of the chunks (see column 9, lines 49-66), and updating the location information by periodically communicating with the servers to obtain changes to the location information (see column 2, lines 50-58 – probes the devices).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to utilize the data collection programs of Staveley et al as an additional program to the programs of the controller mentioned by Bolosky et al. One would have been motivated to do so in order to improve the ability to locate the data files (Bolosky et al: see column 1, lines 38-50).

Referring to claim 16, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality servers configured to store files as chunks (see column 1, line 66 – column 2, line 3 and column 4, lines 37-44 – the disk stores data in blocks which is

considered to represent a chunk since a block is defined as a unit of data or an amount of physical space allocated on a disk to hold that unit of data) and

a master connected to the servers (see column 3, lines 43-45 and Fig 1, items 22, 24 and 26 – the controller is considered to represent the *master*).

Bolosky et al fail to explicitly teach the further limitations wherein the master is configured to communicate with the servers to authoritatively determine location information of the data, the location information being based on which of the servers store the data, periodically communicate with the servers to obtain changes to the location information, and update the location information based on the changes to the location information. Staveley et al disclose a system for monitoring and managing system assets and asset configurations (see abstract). In particular, Staveley et al disclose the further limitations wherein the master is configured to communicate with the servers to authoritatively determine location information of the data (see column 9, lines 49-66), the location information being based on which of the servers store the data, periodically communicate with the servers to obtain changes to the location information (see column 2, lines 50-58), and update the location information based on the changes to the location information (see column 2, lines 50-58 – stores the updated information).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to utilize the data collection programs of Staveley et al as an additional program to the programs of the controller mentioned by Bolosky et al. One would have been motivated to do so in order to improve the ability to locate the data files (Bolosky et al: see column 1, lines 38-50).

Referring to claim 17, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality of servers configured to store data (see column 3, lines 47-50 and Fig 1 – the data servers are considered to represent the plurality of *servers configured to store data*; item 24 represents the data server which contains item 28 which represents a data disk that stores data) and

a master connected to the servers (see column 3, lines 43-45 and Fig 1, items 22, 24 and 26 – the controller is considered to represent the *master*).

Bolosky et al fail to explicitly teach the further limitations wherein the master is configured to communicate with the servers to authoritatively determine location information of the data, the location information being based on which of the servers store the data, instruct one of the servers to perform an action concerning the data, the action causing a change in the location information, and update the location information based on the change to the location information upon completion of the action. Staveley et al disclose a system for monitoring and managing system assets and asset configurations (see abstract). In particular, Staveley et al disclose the further limitations wherein the master is configured to communicate with the servers to authoritatively determine location information of the data, the location information being based on which of the servers store the data (see column 9, lines 49-66), instruct one of the servers to perform an action concerning the data, the action causing a change in the location information (see column 2, lines 50-58), and update the location information

based on the change to the location information upon completion of the action (see column 2, lines 50-58 – updates the information).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to utilize the data collection programs of Staveley et al as an additional program to the programs of the controller mentioned by Bolosky et al. One would have been motivated to do so in order to improve the ability to locate the data files (Bolosky et al: see column 1, lines 38-50).

Referring to claim 18, Bolosky et al disclose a file system (see abstract and Fig 1, item 20), comprising:

a plurality of servers configured to store data (see column 3, lines 47-50 and Fig 1 – the data servers are considered to represent the plurality of *servers configured to store data*; item 24 represents the data server which contains item 28 which represents a data disk that stores data) and

a master connected to the servers (see column 3, lines 43-45 and Fig 1, items 22, 24 and 26 – the controller is considered to represent the *master*).

Bolosky et al fail to explicitly teach the further limitations wherein the master is configured to communicate with the servers to authoritatively determine information regarding the data, instruct one of the servers to perform an action concerning the data, the action causing a state change associated with the information and update state information based on the state change upon completion of the action. Staveley et al disclose a system for monitoring and managing system assets and asset configurations (see abstract). In particular, Staveley et al disclose the further limitations wherein the

Art Unit: 2167

master is configured to communicate with the servers to authoritatively determine information regarding the data (see column 9, lines 49-66), instruct one of the servers to perform an action concerning the data, the action causing a state change associated with the information (see column 2, lines 50-58), and update state information based on the state change upon completion of the action (see column 2, lines 50-58 – updates information).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to utilize the data collection programs of Staveley et al as an additional program to the programs of the controller mentioned by Bolosky et al. One would have been motivated to do so in order to improve the ability to locate the data files (Bolosky et al: see column 1, lines 38-50).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims 8-9 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent No. 5,991,804 issued to Bolosky et al in view of US Patent No 6,973,491 to Staveley et al as applied to claim 7 above, and further in view of US Patent No 6,725,392 issued to Frey et al (hereafter Frey et al).

Referring to claim 8, Bolosky/Staveley discloses a system for providing a master that is further configured to monitor the state of the servers as taught above. However, Bolosky does not explicitly teach the further limitation of the master being configured to exchange heartbeat signals with the servers to determine the state of the servers. Frey et al disclose a system similar to that of Bolosky/Staveley, including the master being configured to exchange heartbeat signals with the servers to determine the state of the servers. In particular, Frey et al teach a system similar to that of claim 7, wherein the master is configured to exchange heartbeat signals with the servers to determine the state of the servers (see column 14, lines 15-19 – the state is whether or not the server is operational).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to add Frey et al's use of heartbeat signals to Bolosky/Staveley's system for monitoring the state of the servers. One would have been motivated to do so since both Frey et al and Bolosky/Staveley deal with fault recovery systems for distributed file systems having a controller and also since the addition of heartbeat signals would enhance the ability to monitor the systems (Frey et al: see abstract; Bolosky et al: see abstract).

Referring to claim 9, the combination of Bolosky/Staveley and Frey et al discloses the system of claim 8, wherein the heartbeat signals include space utilization information (see column 14, lines 15-29 – memory access request is considered to represent space utilization information).

6. Claims 10-11 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent No. 5,991,804 issued to Bolosky et al in view of US Patent No 6,973,491 to Staveley et al as applied to claim 7 above, and further in view of US PGPub 2003/0023898 issued to Jacobs et al (hereafter Jacobs et al).

Referring to claim 10, Bolosky/Staveley discloses a system for providing a master that is further configured to monitor the state of the servers as taught above. However, Bolosky/Staveley does not explicitly teach the further limitation wherein the state of the servers includes information regarding the data stored by the servers. Jacobs et al disclose a system similar to that of Bolosky/Staveley, including the state of the servers including information regarding the data stored by the servers. In particular, Jacobs et al teach a system similar to that of claim 7, wherein the state of the servers includes information regarding the data stored by the servers (see paragraph [0009], lines 6-10 – the version number).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to include Jacobs et al's information regarding the data stored by the servers with Bolosky et al's system for monitoring the state of the servers. One would have been motivated to do so since both Jacobs et al and Bolosky et al deal with sending versions of data from a master server to a second server (Frey et al: see abstract; Bolosky et al: see abstract).

Referring to claim 11, the combination of Bolosky/Staveley and Jacobs et al discloses the system of claim 10, wherein the information includes version numbers of the data (Jacobs et al: see paragraph [0009], lines 6-10).

Response to Arguments

7. Applicant's arguments with respect to claims 1-18 have been considered but are moot in view of the new ground(s) of rejection.

Conclusion

8. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

- US PGPub 2004/0255000 to Simionescu et al titled "Remotely Controlled Failsafe Boot Mechanism and Remote Manager for a Network Device" which focuses on restarting a master computer
- US PGPub 2002/0120763 to Miloushev et al titled "File Switch and Switched File System"

Contact Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Kimberly Lovel whose telephone number is (571) 272-2750. The examiner can normally be reached on 8:00 - 4:30.

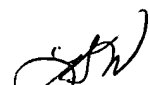
If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John Cottingham can be reached on (571) 272-7079. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Kimberly Lovel
Examiner
Art Unit 2167

kml
5 July 2006


JOHN COTTINGHAM
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100


7 July 2006